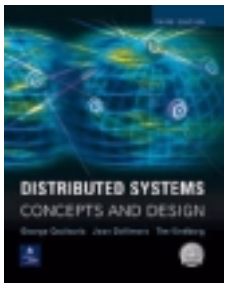

Distributed Systems: Concepts and Design

Edition 3

By George Coulouris, Jean Dollimore and Tim Kindberg
Addison-Wesley, ©Pearson Education 2001

Errata List

As with virtually all books, some bugs were discovered after printing. These errors will be corrected in subsequent printings (impressions). The corrections recorded to date are listed below.

First, second, third, fourth, fifth, sixth and seventh impressions

Page	*existing version*	*correction*	*date*
33	line 8	replace "Common" by "Component"	12-9-00
49	line -2	replace first sentence by "The delay between the start of a message's transmission from one process and the beginning of its receipt by another is referred to as *latency*"	19-2-01
66	lines 12 to14	replace "The Internet is constructed from many *subnets* employing a variety of network technologies. A subnet is a set of interconnected nodes, all of which employ the same technology to communicate amongst themselves." with "The Internet is constructed from many *subnets*. A subnet is a unit of routing (delivering data from one part of the Internet to another); it is a collection of nodes that can all be reached on the same physical network."	21-02-01
67	lines 16 to18	replace with "*Latency* is the delay that occurs after a send operation is executed before data starts to arrive at the destination computer. It can be measured as the time required to transfer an empty message. Here we are considering only network latency, which forms a part of the process-to-process latency defined in Section 2.3.1."	21-2-01
68	line -1	delete "timing"	19-2-01

93	Figure 3.15	Class E addresses need no have a 0 in the fifth binary digit. Delete the box containing a 0 in the row for Class E addresses and replace 27 by 28 in the arrow label. 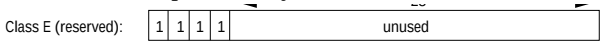	15-03-02
94	Figure 3.16	bottom right cell, replace “128” by “240” and replace “247” by “255”	15-01-01
94	line 5	replace “set to 1” by “set to 0”	26-02-03
96	lines 8-12	replace with: correspondence between IP addresses and addresses for the hosts on the local network and is done using an address resolution protocol (ARP). We now outline the implementation of an ARP for Ethernets. It uses dynamic enquiries in order to operate correctly when computers are added to a local network but exploits caching to minimize enquiry messages.	19-2-01
98	line 16	replace “195.203.255.255” by “203.255.255.255”	9-10-00
108	line 4	replace “model” by “modem”	20-2-03
128	line 29	replace “receiver” by “ receiver (multicast ports are an exception, see Section 4.5.1)	19-2-01
130	line 19	replace “ Any process needing to send or receive messages must first create a socket bound to an Internet address and local port” by “ To send or receive messages a process must first create a socket bound to an Internet address of the local host and a local port.”	19-2-01
130	line -10	replace first sentence by “ Sockets normally provide non-blocking *sends* and blocking *receives* for datagram communication(a non-blocking *receive* is an option in some implementations).”	19-2-01
132	Figure 4.3	*aSocket.close* has been moved to a *finally* clause at the end of the *main* function. The corrected version of this program is available in the file *UDPClient.java* on this website under Additional Material for Chapter 4.	19-2-01
133	Figure 4.4	*aSocket.close* has been moved to a *finally* clause at the end of the *main* function. The corrected version of this program is available in the file *UDPServer.java* on this website under Additional Material for Chapter 4.	19-2-01
136	Figure 4.5	*aSocket.close* has been moved to a *finally* clause at the end of the *main* function. The corrected version of this program is available in the file *TCPClient.java* on this website under Additional Material for Chapter 4.	19-2-01
137	Figure 4.6	*aSocket.close* has been moved to a *finally* clause at the end of the *main* function. The corrected version of this program is availablein the file *TCPServer.java* on this website under Additional Material for Chapter 4.	19-2-01

141	line -5	insert “unsigned” before “long year”	9-10-04
152	Figure 4.15	replace “URL” by “URL or pathname”	18-9-00
152	last line	add “Proxies need the whole URL as shown in the figure, but it is advantageous to send only the pathname in the case of an origin server because there is less to send.”	18-9-00
152	Figure 4.15	replace “//www.dcs.qmw.ac.uk/index.html” by “http://www.dcs.qmw.ac.uk/index.html”	2-3-02
154	line -10	replace “Figure 3.16” by “Figure 3.15”	13-11-00
156	Figure 4.17	*aSocket.close* has been moved to a *finally* clause at the end of the *main* function. The corrected version of this program is available in the file *MulticastPeer.java* on this website under Additional Material for Chapter 4.	19-2-01
169	line 2	replace “Section 4.<<CDR>” by “Section 4.3.1”	15-01-01
169	line 12	replace “Common” by “Component”	12-9-00
205	last line	replace “Exercise 5.13” by “Exercise 5.15”	27-11-01
209	line -16	replace “Section 6.2” by “Section 6.3”	8-2-04
218	line 24	replace first two sentences of paragraph 4 by: “*Migratory* load-sharing systems can shift load at any time, not just when a new process is created. They use a mechanism called *process migration*: the transfer of an executing process from one node to another.”	15-10-01
274	line 5	replace “idempotent” by “self-inverse”	17-01-02
291	line -10	replace “[Steiner et al. 1988]” by “[Neumann and T’so 1994]”	21-9-00
293	line -6	Insert after “Kerberos Version 5” “[Neumann and T’so 1994]	15-9-00
321	line -6	replace “There are three operations for altering directories: *AddName, ReName* and *UnName*.” by “There are two operations for altering directories: *AddName* and *UnName*.”	8-11-00
326	Figure 8.9	replace *link*(*newdirfh, newname, dirfh, name*)→ *status* Creates an entry *newname* in the directory *newdirfh* which refers to file *name* in the directory *dirfh*. with *link*(*newdirfh, newname, fh*)→ *status* Creates an entry *newname* in the directory *newdirfh* which refers to the file or directory *fh*.	02-04-03

328	lines 14-16	replace The file handle returned in the previous step is used as a parameter in the next *lookup* step; the file system identifier in the file handle is first compared with the entries in the remote mount table held in the client to see whether another remote-mounted file store should be accessed. with The file handle returned in the previous step is used as a parameter in the next lookup step. Since file handles are opaque to NFS client code, the virtual file system is responsible for resolving file handles to a local or a remote directory and performing the necessary indirection when it references a local mount point.	09-04-03
330	line -8	insert at start of paragraph There is one value of Tm_{server} for all the data blocks in a file and another for the file attributes.	09-04-03
341	line 9	replace “For each file with a valid token, Venus must send a cache validation request ...” by “Before the first use of each cached file or directory after a restart, Venus therefore generates a cache validation request ... “	8-11-01
366	line 10	replace “IP” by “IN”	19-4-01
387	line 16	replace “17” by “16”	8-11-01
417	Exercise 10.11	replace “Vi[j]” by “Vj[i]”	12-9-00
429	line -3	replace “to all other K-1 members of Vi" by "all K members of Vi (including itself)”	8-2-04
429	line -2	replace "all K-1 reply messages" by "all K reply messages"	8-2-04
430	Fig 11.6 line 6	delete "- {pi}"	8-2-04
430	Fig 11.6 line 7	replace "(K-1)" by "K"	8-2-04
430	Fig 11.6 line 9	delete "(i != j)"	8-2-04
430	Fig 11.6 line 19	delete "- {pi}"	8-2-04
430	Fig 11.6 line 20	delete "(i != j)"	8-2-04
430	line -8	replace "all the other K - 1 members of Vi" bye "all K members of Vi (including itself)"	8-2-04
431	lines 1-2	replace “then it is possible for p_1 to reply to p_2 but hold off p_3; for p_2 to reply to p_3 but hold off p_1; and for p_3 to reply to p_1 and hold off p_2” by “it is possible for p_1 to reply to itself and hold off p_2; for p_2 to reply to itself and hold off p_3; for p_3 to reply to itself and hold off p_1”	16-4-04
440	penultimate line	replace “protocol” by “description”	14-9-00

441	last paragraph, sentences 2 and 3	replace by: "The validity property holds because IP multicast has that property. For agreement we require, first, that a process can always detect missing messages. That in turn means that it will always receive a further message that enables it to detect the omission. As this simplified protocol stands, we guarantee detection of missing messages only in the case where correct processes multicast messages indefinitely. "	19-10-00
441	last paragraph sentence 4	replace "The agreement property holds as long as" by " Second, the agreement property requires that"	12-10-00
442	line 1	replace first sentence by: "Neither of the assumptions we made to ensure agreement is practical (see Exercise 11.14)."	12-10-00
442	line 2	replace " However, validity and agreement are practically addressed in the protocols" by: "However, agreement is practically addressed in the protocols	12-10-00
446	Figure 11.14	replace *On B-deliver(<"order", i, S>) with g = group(m)* *wait until <m, i> in hold-back queue and S = r_g+1;* *TO-deliver m; // (after deleting it from the hold-back queue)* *r_g= S;* by *On B-deliver(m_{order}=<"order", i, S>) with g = group(m_{order})* *wait until <m, i> in hold-back queue and S = r_g;* *TO-deliver m; //after deleting it from the hold-back queue)* *r_g= S+1*	20-03-01
449	Figure 11.16	replace "*On B-deliver*(<V^g_j, *m*>) *from* p_j *, with g = group*(*m*)" by "*On B-deliver*(<V^g_j, *m*>) *from* p_j *(j≠i), with g = group*(*m*)*"*	10-04-03
454	line -10	Insert "For the case where a majority of processes are correct," before "we construct a solution to C from IC...'.	23-07-02
455	line 4	Insert "In systems with crash failures," before beginning of paragraph "Solving consensus is equivalent...".	23-07-02
457	line 3	replace "Figure 11.18" by "Figure 11.19"	27-10-02
463	Exercise 11.13	replace first sentence by "Explain whether the algorithm for reliable multicast over IP works for open as well as closed groups"	14-9-00
481	line 19	Insert "(and their descendants)" after "1. Subtransactions at one level "	15-04-04
482	line 9	replace "Transfer $100 from A to B" by "Transfer $100 from B to A"	16-8-02

498	line 22	replace "T_v" by "T_j"	12-11-01
524	paragraph 2	replace with "When a subtransaction completes, it makes an independent decision either to commit provisionally or to abort. A provisional commit is different from being prepared to commit: it is not backed up on permanent storage. If the server crashes subsequently, its replacement will not be able to commit. After all subtransactions have completed, the provisionally committed ones participate in a two-phase commit protocol, in which servers of provisionally committed subtransactions express their intention to commit – or abort if an ancestor has aborted. A prepared commit guarantees a subtransaction will be able to commit, whereas a provisional commit only means it has finished correctly."	15-04-04
551	Exercise 13.9 line 3	replace " Figure 13.6" by "Exercise 13.4"	13-7-00
562	lines -5 to -11	replace paragraphs on Agreement and Integrity with: *Agreement*: Correct processes deliver the same sequence of views (starting from the view in which they join the group), and the same set of messages in any given view. That is, if a correct process delivers message *m* in view *v(g)*, then all other correct processes that deliver *m* also do so in the view *v(g)*. *Integrity*: If a correct process *p* delivers message *m*, then it will not deliver *m* again. Furthermore, $p \in group(m)$ and the process that sent *m* is in the view in which *p* delivers *m*.	26-2-03
639	Figure 16.2 line 2	replace "inta" with "int a"	14-9-00
644	line 23	replace "Chapter 11" with "Chapter 10"	14-9-00
670	line -3	replace "Seetharamanan" by "Seetharaman"	19-7-02
673	Section 17.2.1	This example is described for Java 2 version 1.3 or earlier. See 'Java CORBA with Java 2 version 1.4' under 'Additional Material' on *www.cdk3.net*. (under Supplementary Material for Chapter 17.)	2-4-03
683	Figure 17.8 lines -5 and -2	add "{" and replace "strings" by "string s"	28-11-00
698	Exercise 17.19	replace "17.13" by "Exercise 17.18"	13-7-00
714	line 14	insert the following: "a 'copy-inherited' region is one that is" after "*copied:*"	6-9-01
744	Insert additional reference	Mockapetris 1987," Mockapetris, P. (1987).*Domain names - concepts and facilities*. Internet RFC 1034.November.	15-4-04

745	Insert additional reference	Neumann and Ts'o 1994 Neuman, B.C. and Ts'o, T. 1994. Kerberos: An Authentication Service for Computer Networks, *IEEE Communications*, vol. 32, no. 9, pp. 33-38. Sept. 1994. http://nii.isi.edu/publications/kerberos-neuman-tso.html	15-9-00
751	penultimate reference	replace "Spasojevic R"by "Spasojevic M"	2-8-00
751	fifth reference	replace "Seetharamanan" by "Seetharaman" (twice)	19-7-02
752	fourth reference	replace "Srinivasa" by "Srinivasan"	15-04-04
759	Under CORBA	replace "interoperable Inter-ORB protocol (IIOP)" by "Internet Inter-ORB protocol (IIOP)"	15-4-04

Most recently reported error 9 October 2004